Fermilab

New Multithreaded Code for Calculating Longitudinal Collective Instabilites using Computers with Multiprocessors

Cheng-Yang Tan

Beams Division/Tevatron

ABSTRACT: We have developed a new mulithreaded code with pthreads for calculating longitudinal collective instabilites on computers with multiprocessors. We have selected pthreads as the basis for multiprocessing because it is portable, as such we are able to port this code to Solaris, IRIX and OS/2 platforms. We will demonstrate that when there are four cavities and 36 bunches in the simulation, our code shows a speed increase of $> 3\times$ compared to single processor code when run on a symmetric multiprocessing (smp) machine.

INTRODUCTION

There are many well established codes used in longitudinal simulations of beam instabilities, like ESME[1], which run on a single processor. With the introduction of computers with many processors (multiprocessors) like SUN Sparcs with four or more processors, it seemed to be a waste not to use the other processors in our simulations. Modern workstations with many processors are usually symmetric multiprocessing[†] and the vendors have written their operating systems to support this architecture. This means that the vendor has already written the appropriate libraries which allow the programmer to easily write programmes which can run on multiprocessors. The usual challenge which faces the physicist is to come up with a parallel processing algorithm which takes advantage of the hardware. We will show in this paper that for longitudinal simulations with many RF cavities, the parallel algorithm which we will choose is very natural and indeed takes advantage of the multiprocessor machine. The next challenge which we will face is the choice of libraries and language. The choice for us will again obvious: we will choose a library which is cross-platform and an object-oriented language.

We will, in the rest of the paper, introduce the idea of threads and the pipeline model used in the simulation. Plus results we will use to check that our results are correct when compared with established theory. And of course, the speed increases compared with single processor code.

[†] Symmetric multiprocessing (smp) means that all CPUs are treated the same, as opposed to asymetric multiprocessing where some CPUs are special and others which are slaves.

THREADS

A thread is defined as a sequence of instructions to be executed within a programme. Most programmes are single threaded code which consists of one thread of execution which starts in main(). Before the invention of threads, the UNIX way of doing things in parallel is to use the fork/exec model. This model spawned several processes each a thread of execution which can then be run in parallel. However, each new process that is created by fork/exec require the operating system to make a child process from its parent process by copying over the instruction, user-data and system-data segments of the parent and then executing the child process as well as the parent process. This method is expensive because of the overhead required for each process creation. Furthermore, communication between parent and child need external channels like pipes, sockets, memory maps (mmaps) etc. because resources held by parent and child are private and neither parent or child can peek into each others resources without using these mechanisms. The fork/exec model can therefore be thought of as many processes each having its own thread of execution.

Contrast this with a multithreaded programme where there is one process with many threads of execution. This means that different parts of the same code can be executing in parallel. The advantages over fork/exec are immediate : all threads share the same resources and the overhead of creating a new process is eliminated. The downside is that because all threads share the same resources, there must be software mechanisms, like locks and semaphores, which preserve the integrity of the data and prevent race and deadlock conditions. For example, suppose there are two threads A and B which use a piece of shared data. At some point in the process, thread A updates the data, so thread A must prevent thread B from reading the data before thread A completes the update, i.e. prevention of a race condition. Furthermore, if a situation arises where thread A cannot finish the update unless thread B reads the data then we have a deadlock condition. This means that our

application is stuck in a never ending wait between threads A and B. To prevent these type of problems, judicious use of locks and semaphores is essential. Unfortunately, this adds to the overhead of multithreaded code when compared with single threaded code which means that there are scenarios where single threaded code will run faster than multithreaded code.

Pthreads

POSIX threads or pthreads is a cross-platform implementation of threads whose programming interface is specified by IEEE POSIX 1003.1c standard (1995). This standardization allow us to port the code with minimal changes to two distinct UNIX platforms: Solaris, IRIX and one PC platform: OS/2. Although pthreads is highly portable, there are some shortcomings. For example, a feature which we have thought will be useful is the explicit specification of the mapping between thread and CPU. This feature does not exist in pthreads and thus under this programming environment, the programmer defers the dispatch of threads to CPUs to the operating system. If the programmer really wants to force a thread to a CPU, (s)he must use non-portable system calls.

MODEL

We will describe here the pipeline model used to parallelize the longitudinal simulation code. The pipeline model, which is extremely similar to the Ford assembly line, is shown in Figure 1. We start first with a computer with M processors. For illustration, let us suppose that in our simulations we have $1 < n \leq M$ RF cavities. We can naturally place a single thread of RF cavity code on each CPU as shown in Figure 1. We suppose that we have N bunches numbered bunch0, bunch1,..., bunch$(N-1)$ for the simulation. At the start of the simulation, we put bunch0 into RF Cavity 0 with the other bunches waiting to be processed. Once RF Cavity 0 is done with bunch0, it is sent to RF Cavity 1 and bunch1 is sent into RF Cavity 0. Both RF Cavity 0 and RF Cavity 1 can then work on their bunches in parallel. Once RF Cavities 0 and 1 are done with their respective bunches, bunch0 goes to RF Cavity 2, bunch1 goes into RF Cavity 1 and bunch2 goes into RF Cavity 0 where they are again processed in parallel. As this process is repeated for the remaining bunches, we can see that eventually all the processors will be working in parallel. The RF cavities themselves form a pipeline and the bunches propagate through this pipeline starting from RF Cavity 0 and move towards RF Cavity$(n-1)$ which is at the end of the pipeline. The first bunch to reach the end is bunch0 which means that bunch0 has completed one turn through the accelerator and is ready to be fed back into the beginning of the pipeline at RF Cavity 0 again. This is done *ad infinitum* for all the bunches until the required number of turns have been met.

The advantage of this model is that when there is more than one CPU in the computer and more than one RF cavity in the simulation we can efficiently utilize the computer resources to speed up the calculations by placing a thread on each CPU. However, when there is only one RF cavity in the simulation, the overhead used to set up and maintain the pipeline, which basically consists of one thread only, will be much greater than sin-
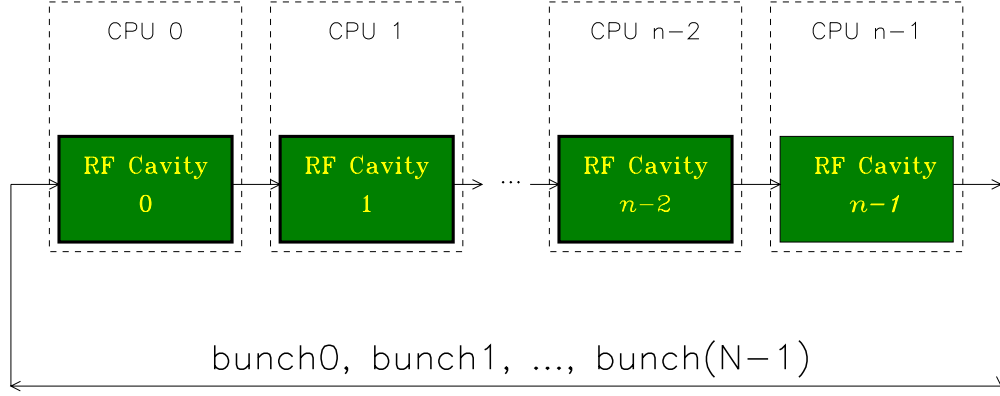
Pipeline Model for SMP

Figure 1   This figure shows the pipeline model used in the simulation. Each RF cavity is placed on a CPU and each bunch is evolved through each cavity like a pipeline.

gle threaded code. This means that the simulation will run slower using multithreaded code compared with single threaded code. Similarly, for the case when there is only one CPU and many RF cavities in the simulation, i.e. many threads and one CPU, there is also a comparable slowdown.

It is also important to keep the pipeline from stalling. Consider the case when there is only one bunch and many RF cavities. Then all the RF cavities except one are sitting idle in the simulation, thus again the thread overhead overwhelms any advantage of using multiprocessors. We can always tell that the pipeline is stalled because when the number of bunches is increased with the number of RF cavities fixed, the speed of the simulation will increase and tend to an asymptotic limit which is when all the CPUs in the pipeline

are working with full steam.

Finally, it is also not too much of a stretch of the imagination to think of each RF cavity as an object which is mapped to each CPU and each bunch is an object which is manipulated by each RF cavity. Thus the natural language to code the simulation in is in some object oriented language like $C++$.

The equations of motion for longitudinal macro-particle tracking have been derived in many places. We quote from Edwards the equations of motion when there is only one RF cavity in the system. These equations can be easily be generalized to many RF cavities

$$\left.\begin{array}{c} \phi_{n+1} = \phi_n + \eta\omega_{\rm rf}\tau_{n+1}\left(\dfrac{\Delta p}{p}\right)_{n+1} \\[2mm] E_n = E_n + eV\sin\phi_n \end{array}\right\} \tag{1}$$

where

ϕ_n is the phase accumulated by the particle after going through one turn.

η is the slip factor.

$\omega_{\rm rf}$ is the RF frequency.

τ_n is the time taken by the particle to go through the nth turn.

$(\Delta p/p)_n$ is the fractional change in momentum between the particle and the synchrononous particle in the nth turn.

E_n is the energy of the particle in the nth turn.

V is the voltage of the cavity.

Besides (1), we also need to know how the resonant modes of the cavity kick the bunch as it goes through the cavity. This is calculated using the wake function $W_0'(z)$ quoted from Chao

$$W_0'(z) = \begin{cases} 0 & \text{if } z > 0, \\[2mm] \alpha R_s & \text{if } z = 0, \\[2mm] 2\alpha R_s e^{\alpha z/c}\left(\cos\dfrac{\bar{\omega}z}{c} + \dfrac{\alpha}{\bar{\omega}}\sin\dfrac{\bar{\omega}z}{c}\right) & \text{if } z < 0. \end{cases} \tag{2}$$

where $\alpha = \omega_r/2Q$, and $\bar{\omega} = \sqrt{\omega_r^2 - \alpha^2}$, R_s is the shunt impedance and Q is the quality factor of the resonance. $z = s - ct$ is the position of the particle w.r.t. reference particle.

These two set of equations form the backbone of how the particles in the bunches evolve as they go through the accelerator. To verify that the solution from the simulation

at least agrees with simple cases where there are analytic solutions, we use the solutions from Sacherer who has calculated the growth rate for coupled bunch modes. The formulæ which we quote below is a particular form of Sacherer's formula which is derived by Ng.

Suppose we have a phase space $(x, \Delta E)$ distribution ψ of a bunch which can be separated into a stationary part ψ_0 and a perturbed part ψ_1 i.e.

$$\psi(x, \Delta E; s) = \psi_0(x, \Delta E) + \psi_1(x, \Delta E; s) \tag{3}$$

where s is the position of the synchronous particle. Then if we assume that the solution is sinusoidal in time, then Sacherer showed that

$$\text{growth of } \psi_1 \sim e^{t/\tau_{m,\mu}} \tag{4}$$

where $1/\tau_{m,\mu}$ is the growth rate of the μ'th coupled bunch mode from the mth pole phase space oscillation. In particular, the formula for the growth rate due to dipole modes to first order for M *equally* spaced bunches is

$$1/\tau_{1,\mu} = \frac{\eta e^2 N M \omega_r}{2\beta^2 E_0 \tau_0^2 \omega_s} \left[\text{Re} Z_0^{\|}(qM\omega_0 + \mu\omega + \omega_s) - \text{Re} Z_0^{\|}(q'M\omega_0 - \mu\omega + \omega_s) \right] \tag{5}$$

where

$0 \leq \mu < M$ is the μ'th coupled bunch mode.

N is the number of particles in a bunch.

M is the number of bunches.

q is the revolution harmonic.

ω_0 is the revolution frequency.

ω_r is the frequency of the resonance.

ω_s is the synchrotron frequency.

τ_0 is the revolution period.

E_0 is the energy of the bunch.

and $Z_0^{\parallel}$ is the longitudinal impedance of the mode. For sharp resonances which can be approximated with an RLC resonator, $Z_0^{\parallel}$ can be modeled as

$$Z_0^{\parallel}(\omega) = \frac{R_s}{1 + iQ\left(\frac{\omega_r}{\omega} - \frac{\omega}{\omega_r}\right)} \tag{6}$$

where R_s is the shunt impedance and Q is the loaded quality factor of the resonance.

It must be emphasised that the theory is only valid for equally spaced bunches. When the bunches are not equally spaced, the theory is still good at giving qualitative growth or damping results.

In particular when $\mu = 0$, we can see that the sign of $1/\tau_{1,0}$ is dependent on which side the resonance ω_r is w.r.t. the revolution harmonic $qM\omega_0$. See Figure 2. Above transition, $\eta > 0$ and if $\mathrm{Re}Z_0^{\parallel}$ at the upper synchrotron sideband is larger than at the lower synchrotron sideband, the system is unstable because $1/\tau_{1,0} > 0$. The system will be stable if $1/\tau_{1,0} < 0$ when $\mathrm{Re}Z_0^{\parallel}$ at the upper synchrotron sideband is smaller than at the lower synchrotron sideband.
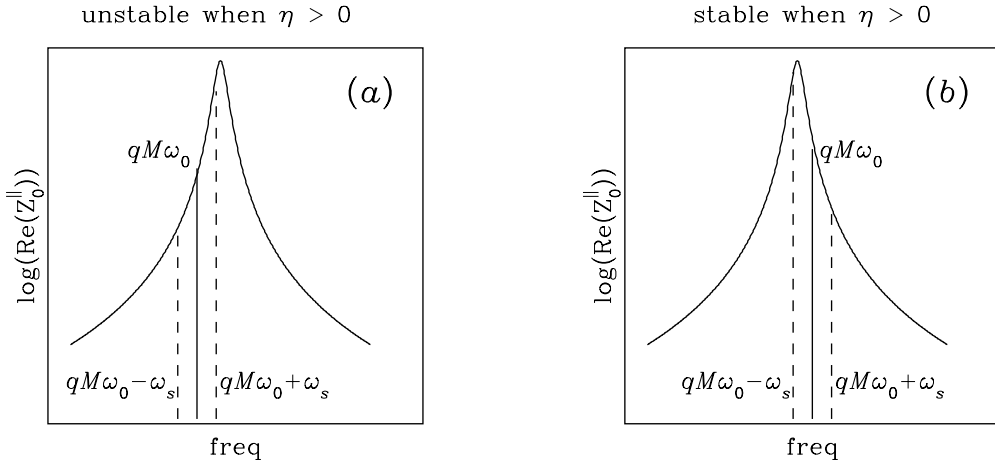

Figure 2 This figure shows how the position of the resonance w.r.t. the revolution harmonic $qM\omega_0$ affects the sign of $1/\tau_{1,0}$. Above transition (a) is unstable and (b) is stable in this figure.

10

Check

Using the results shown above, we checked our simulation with (5) for the case when we have three equally spaced point bunches in the simulation. The resonances and other parameters used in the simulation are shown in Table 1. We can see from Figure 3 that theoretical damping and growth rates are very close to the results of the simulation for short time scales. As time gets longer, the theory and simulation start disagreeing. This should not be too surprising because Sacherer's theory only gives a first order result plus it is only valid at the *start* of the instability.

Table 1. Parameters used in the simulation

Symbol	Description	Value
μ	coupled bunch mode	0
N	# particles in one bunch	2.7×10^{11}
M	# of bunches	3
q	revolution harmonic #	2174
E_0	Energy of the bunch	150 GeV
ω_0	revolution frequency	$(2\pi \times 47.7125 \times 10^3)$ s^{-1}
ω_s	synchrotron frequency	$(2\pi \times 88.296)$ s^{-1}
ω_r	frequency of the stable resonance	311.180 MHz
	frequency of the unstable resonance	311.183 MHz
R_s	shunt impedance of the resonance	1.73 MΩ
Q	loaded quality factor of the resonance	134×10^3

Using the values in Table 1, and substituting them into (5) we find that for the stable resonance, $1/\tau_{1,0} = -1.573$ s^{-1}, while for the unstable resonance $1/\tau_{1,0} = +0.537$ s^{-1}.
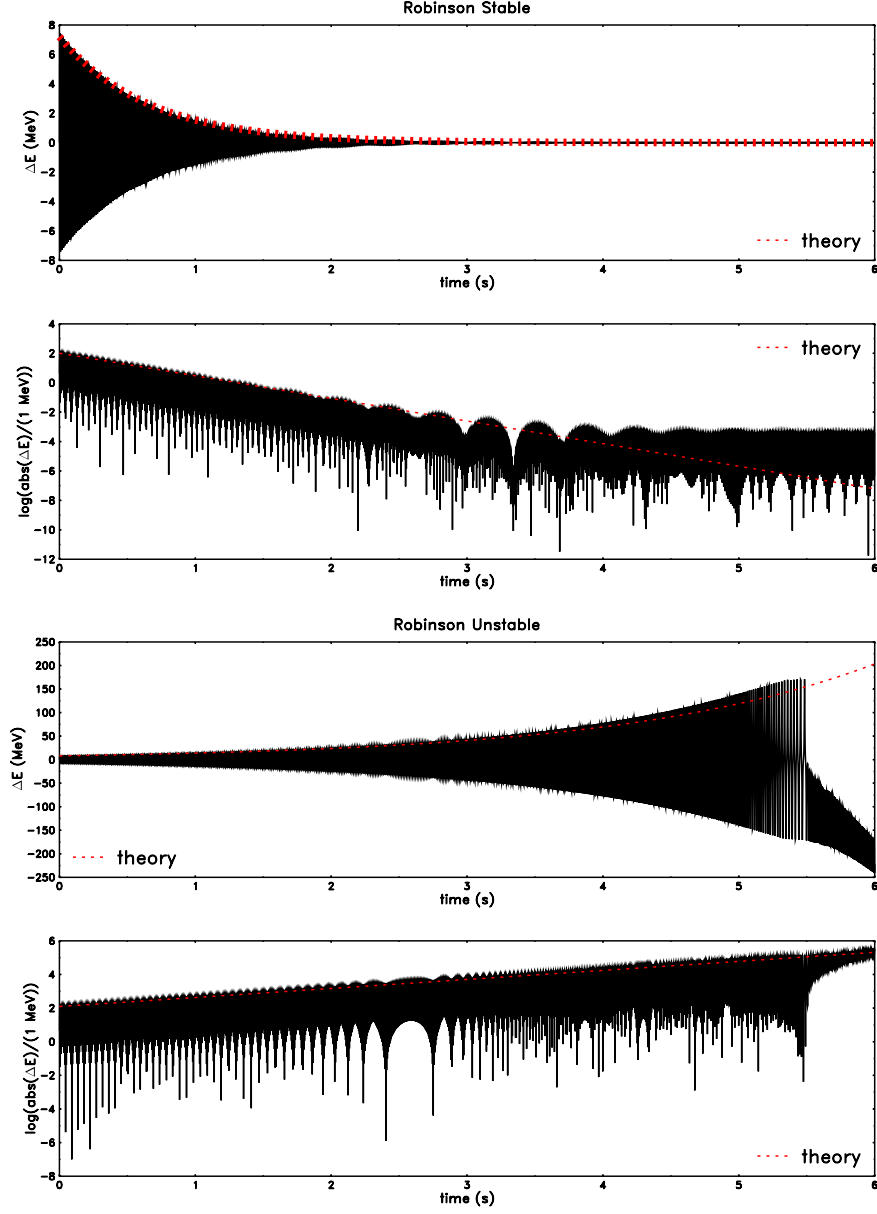
Figure 3 The simulation results and the theoretical growth rates are shown here. The top two graphs show the case when we have Robinson stability, and the bottom two graphs when we have Robinson instability (the particle leaves the bucket at about 5.5s).

SPEED RESULTS

For amusement we show the results of *top* in Figure 4. This shows that our programme *bl* is using 330% more CPU than the next highest user thus leaving him in the dust! For a definitive check of the speedups, we compare completion times with single threaded code and multithreaded code and saw $> 3\times$ increase in speed when we have four RF cavities and 36 bunches in the simulation. Table 2 shows the actual speed increase with timed results on a Sparc workstation with six processors (unfortunately we can only use four, see *CONCLUSION* for the reason) and running SunOS 5.6. We define the speedup factor to mean

$$\text{speedup factor} = \frac{\text{completion wall-clock time with single thread}}{\text{completion wall-clock time with multithreads}} \tag{7}$$

Notice from Table 2 that there is a cross-over point when a single threaded simulation is *faster* than a multithreaded simulation when we have two RF cavities and < 4 bunches in the simulation. The single threaded simulation is faster than the multithreaded simulation because the pipeline is stalled from too few bunches.

The other extreme case is when we have more RF cavities than CPUs, which is illustrated with the eight RF cavities case, and even then, the multithreaded performance is still much better than single threaded performance.

Table 2. Comparison between single and multithreaded performance

# RF cavities	# CPUs used	# bunches	speedup factor
8	4	36	3.2×
4	4	36	3.2×
2	2	36	1.7×
2	2	8	1.2×
2	2	4	1.0×
2	2	2	0.5×

Figure 4   This shows the output of *top* which is used to monitor CPU usage. As can be seen from the line containing USERNAME cytan, our programme *bl* is using 332.8% CPU while other users are using < 100% CPU. This particular machine has 16 CPUs and we are using 4 RF cavities and 36 bunches in the simulation.

CONCLUSION

As can be seen from Table 2, the speedup factor is dependent on the number of RF cavities as well as the number of bunches. Thus to determine whether there is an advantage in doing a simulation with multiprocessors, we must first do a computation with a short simulation time to determine which method is better before committing ourselves. Furthermore, not only do we have to consider the non-linear scaling between the speedup factor and the number of CPUs, there are also hard and soft speed limits which prevent us from going infinitely fast

(*i*) A hardware limit which is the overhead in communications between processors. Even when we have four available processors on the Sparc to simulate four RF cavities, we do not get $4\times$ increase in speed.

(*ii*) A software limit put in by the system administrators which prevent us CPU hogs from using more than four processors.

14

Notwithstanding the above two limits, we have shown that multithreaded code which runs on modern smp machines can indeed make our simulations run much faster and clearly in the 21st century this is the way to go.

ACKNOWLEDGEMENTS

The author wishes to thank Jim MacLachlan and King-Yuen Ng for many hours of fruitful discussions.

REFERENCES

[1] *ESME*, J.A. MacLachlan, `http://www-ap.fnal.gov/ESME`.

[2] *An Introduction to the Physics of High Energy Accelerators*, D.A. Edwards and M.J. Syphers, Wiley Series in Beam Physics, 1993.

[3] *Physics of Collective Beam Instabilities in High Energy Accelerators*, A.W. Chao, Wiley Series in Beam Physics, 1993.

[4] *Introduction to Collective Instabilities – Longitudinal and Transverse*, K.Y. Ng, TM-2053.

[5] *Programming with POSIX Threads*, D.Y. Butenhof, Addison Wesley Publishing Company, 1997.